On Verifying and Maintaining Connectivity of Interval Temporal Networks

Eleni C. Akrida[1]([✉]) and Paul G. Spirakis[1,2]

[1] Department of Computer Science, University of Liverpool, Liverpool, UK
Eleni.Akrida@liverpool.ac.uk
[2] Computer Technology Institute and Press "Diophantus" (CTI), Patras, Greece
P.Spirakis@liverpool.ac.uk

Abstract. An interval temporal network is, informally speaking, a network whose links change with time. The term *interval* means that a link may exist for one or more time intervals, called *availability intervals* of the link, after which it does not exist (until, maybe, a further moment in time when it starts being available again). In this model, we consider continuous time and high-speed (instantaneous) information dissemination. An interval temporal network is connected during a period of time $[x, y]$, if it is connected for all time instances $t \in [x, y]$ (instantaneous connectivity). In this work, we study instantaneous connectivity issues of interval temporal networks. We provide a polynomial-time algorithm that answers if a given interval temporal network is connected during a time period. If the network is not connected throughout the given time period, then we also give a polynomial-time algorithm that returns large components of the network that remain connected and remain large during $[x, y]$; the algorithm also considers the components of the network that start as large at time $t = x$ but dis-connect into small components within the time interval $[x, y]$, and answers how long after time $t = x$ these components stay connected and large. Finally, we examine a case of interval temporal networks on tree graphs where the *lifetimes* of links and, thus, the failures in the connectivity of the network are not controlled by us; however, we can "feed" the network with extra edges that may re-connect it into a tree when a failure happens, so that its connectivity is maintained during a time period. We show that we can with high probability maintain the connectivity of the network for a long time period by making these extra edges available for re-connection using a randomised approach. Our approach also saves some cost in the design of availabilities of the edges; here, the cost is the sum, over all extra edges, of the length of their availability-to-reconnect interval.

1 Introduction and Motivation

A great variety of systems in society, technology and nature can be modelled as networks, linked with edges; from the Internet to the web of social acquaintances,

from the transport network of a city to the nervous system of the human body. The structure of a network describes the several connections between the participating entities and helps us understand or predict the behaviour of dynamical systems. However, in many cases the links between the participating entities do not always remain active but change or disappear as time progresses. A *temporal network* is, informally speaking, a network that changes with time. Both traditional and modern networks, such as communication networks, social networks, transportation networks and physical systems, can be modelled as temporal.

Dynamic networks in general have been attracting attention over the past years, exactly because they model real-life applications. The study of temporal networks in particular is quite interdisciplinary, which is also reflected in literature where the object of study may have different names - temporal graphs, temporal networks, evolving graphs, time-stamped graphs etc. Kempe et al. [14] considered the single-labelled discrete-time model of temporal graphs, where every edge may become available (for use) only at a discrete moment in time, called the *label* of the edge; their main motivation was to examine how basic graph properties change in this temporal setting. In their multi-labelled model, Mertzios et al. [17] extended the model of [14] to many labels per edge and mainly examined the number of labels needed for a temporal design of a network to guarantee several graph properties with certainty. They also provided an algorithm to compute foremost time-respecting paths; in this discrete-time model, a time-respecting path is a path in which successive edges have strictly increasing time labels and a foremost time-respecting path is one that *reaches* the destination vertex at the earliest possible time. Random edge availabilities in the discrete-time model of temporal networks were first considered by Akrida et al. [1] in order to study the Expected Temporal Diameter of temporal graphs.

Assuming the availability of an edge for a whole time-interval $[t_1, t_2]$ or multiple such time-intervals, and not just for discrete moments, is a clearly natural assumption since time is indeed a continuous measure. Bui-Xuan et al. [4] consider a class of dynamic networks where the changes in the topology can be predicted in advance and in which each node and each edge comes with a list of time intervals; they give algorithms for computing foremost time-respecting paths, shortest (minimum hop count) time-respecting paths and fastest (minimum time) time-respecting paths in this model. Fleischer and Tardos [11] consider a continuous-time model of dynamic graphs and prove continuous versions of known discrete-time flow algorithms for dynamic flow problems. Fleischer and Skutella [10] also engage in the study of flows in a continuous-time model of dynamic graphs. Further related work includes [2,3,5,6,9,12,13,15,16,18–21].

1.1 Our Contribution

In this work, we restrict our attention to *continuous time* and consider systems in which only the connections between the participating entities may change, while the entities remain unchanged. So we consider networks of a fixed vertex set, each edge e of which is available over a set of time intervals $L_e = \{[t_1, t'_1], \ldots, [t_k, t'_k]\}$. Each interval indicates a period of availability of e; the unprimed times mark

the start of the availability period and the primed times mark the end. This is a model that could naturally represent several systems, such as proximity networks where a link may represent that two entities have been close to each other for some extent of time, or infrastructural systems like the Internet, or even seasonal food webs where a time interval may represent the fact that one species is the main food source of another for a specific period of the year.

We give a polynomial-time deterministic algorithm that decides if a given interval temporal network is connected during a given period (cf. Sect. 3); if the network is not connected, the algorithm returns the maximal interval from the beginning of the given period during which the network stays connected. We then provide a polynomial-time algorithm that decides if a given interval temporal network has *large enough* connected components during a given time period; here, the size of the components in question is determined by a parameter provided by the user as input to the algorithm (cf. Sect. 4). Finally, we provide a probabilistic analysis of a scenario where the *lifetime* of the intervals assigned to the edges of a network on a tree graph are not designed via a deterministic process and are unknown to us; instead, the edges may fail unexpectedly and we are required to supply the network with more available edges so that, when a break in the connectivity of the network happens, we can re-connect it. We wish not to keep all these extra edges available for re-connection at all times, i.e., we wish to maintain connectivity but by paying a low cost on keeping extra edges available. Assuming that the cost of keeping additional edges available is linear to the sum of lengths of their availability intervals, we show a low cost construction. Other work for maintaining some structure or property like connectivity in probabilistic dynamic graphs includes [7,8].

2   Preliminaries

We focus here on networks, the links of which are not always available. The availability of a link is described via a set of time intervals, one set per edge.

Definition 1 (Interval Temporal Network). *Let $G = (V, E)$ be a (di)graph. An interval temporal network on G is an ordered triplet $G(L) = (V, E, L)$, where $L = \{L_e = \{[t_1, t_1'], \ldots, [t_{k_e}, t_{k_e}']\},$ for some $k_e \in \mathbb{N},\ t_i, t_i' \in \mathbb{R}^+, t_i < t_i',\ i = 1, 2, \ldots,\ k_e : e \in E\}$ is an* assignment *of availability intervals to the edges (arcs) of G. L is called a* labelling *of G.*

The availability intervals of an edge (arc) e represent the *continuous time intervals* at which e is active. When we say that an edge (arc) is active or available during the interval $[a, b]$, for some $a, b \in \mathbb{R}$, it means that the edge exists in the network $\forall t \in \mathbb{R}^+,\ t \in [a, b]$. For the analysis throughout the paper, we assume the intervals $[t_1, t_1'], \ldots, [t_{k_e}, t_{k_e}']$ to be disjoint.[1] Every time a change in the network happens, i.e., an edge starts or stops being available, we have changes

---

[1] We can assume this, because if an edge $e \in E$ has overlapping availability intervals, then we can consider their union as an availability interval of e.

in the topology of the network; so, in a sense, an interval temporal network can be viewed as a sequence of graphs, one after every topology change. However, representing such networks as evolving graphs, i.e., the sequence of states of the network after each change, is not as efficient. The interval representation is indeed a very compact representation of such kinds of evolving graphs.

A basic assumption that we follow is that when a message or an entity passes through an available link at time t, then it can pass through a subsequent link only at some time $t' \geq t$ and only at a time at which that link is available. However, unlike what is assumed in the discrete-time model of [1,14,17], here we consider instant information dissemination through a path of the underlying (di)graph, if the consecutive edges (arcs) are consistently labelled. In fact, our model considers very high speed of information dissemination, resembling fibre-optic communication, but the small time needed to send a message through a link is considered negligible for the analysis. Consider, for example, an interval temporal network $G(L)$ and a path p of $G(L)$ such that all edges of p are available at some time $t = t_0$; if some information starts at time t_0 from one endpoint of p, it can arrive at time t_0 to the other endpoint.

Definition 2 (Connectivity of Interval Temporal Networks). *An interval temporal network $G(L) = (V, E, L)$ is connected at a given time instance t_0 if the edges that are available at time t_0, i.e., the edges that have an availability interval which includes t_0, induce a spanning tree.*

3 Connectivity of Interval Temporal Networks During a Given Time Period

A fundamental issue for any given network, dynamic or not, is to verify if the network is connected (over time, in the dynamic case), i.e., information can travel via edges between any ordered pair of vertices in it. In this section, we consider interval temporal networks and address the issue of their connectivity.

One can think of an interval temporal network as a dynamic network, where the changes in the topology of the network happen whenever an availability interval of an edge starts or finishes, but can view it as static in between these (instantaneous) changes. Since information can travel instantaneously in interval temporal networks, for such a network to be connected over a time period, all the instances of the "static" networks that are formed during that period need to be connected.

We provide below a polynomial-time procedure to determine if a given interval temporal network is connected throughout a particular time period. Henceforth, we denote by $E(t)$ the set of edges that are available at time t, and t is *not* the finish time of the availability interval that includes t.[2]

---

[2] $E(t)$ are the edges that are available at t and do not stop being available (immediately) after time t.

Theorem 1. *There is a polynomial-time algorithm (cf. Algorithm 1) which, given an interval temporal network $G(L)$ on n vertices and numbers $x, y \in \mathbb{R}^+$, $x < y$, answers whether $G(L)$ is connected during the time period $[x, y]$, i.e., is connected for every time instance $t \in [x, y]$. If for some $a \in [x, y]$, $[x, a] \subseteq [x, y]$ is the maximal sub-period of $[x, y]$ during which $G(L)$ remains connected, then the algorithm also returns the length of that period, $a - x$.*

Algorithm 1. Connectivity of interval temporal networks

Input: A temporal network $G(L)$ of n vertices and numbers $x, y \in \mathbb{R}^+$ such that $x < y$
Output: Answer if $G(L)$ remains connected during the time interval $[x, y]$

1: **if** $E(x)$ induces a spanning tree, T, of G **then**
2:     Sort the edges in T according to the finish time of their availability interval; //For every edge in T, we only consider the interval that includes x
3:     Let $A = \{e_i, \text{ with interval } [a_i, b_i] : i = 1, \dots, n-1\}$ be the sorted list;
4:     **if** $b_1 \geq y$ **then**
5:         **return** "Network is connected" **and** "Duration of survival =" $y - x$; //If all edges in T remain available until (at least) time y
6:     **else**
7:         $E' := \{e \in E(T) : b_e = b_1\}$;   //$b_1$ is the first time instance at which T becomes disconnected. E' is the set of edges of T that stop being available at time b_1.
8:         $T := T \setminus E'$;   //T is now a forest, i.e., consist of a collection of trees
9:         Remove E' from A;
10:        Let $T_1, T_2, \dots, T_i, i \in \mathbb{N}$ be the connected components of T;
11:        **while** T is disconnected **do**
12:            **if** $\exists j, k = 1, 2, \dots, i : \exists e = (u, v) \in E(b_1) : u \in V(T_j) \wedge v \in V(T_k)$ **then**
13:                Find the T_j, T_k trees of T that e connects;   //If there is an edge of G with endpoints in different connected components of T and is available at time b_1, then add it to T
14:                Merge T_j, T_k and e into a single tree;
15:                Update the number i of connected components of T;
16:                Insert e in the sorted list A;
17:            **else**
18:                **return** "Network is disconnected" **and** "Duration of survival =" $b_1 - x$;

19:                Break;
20:        Go to line 4
21: **else**
22:     **return** "Network is disconnected" **and** "Duration of survival =" $t - x$;

Description of the Algorithm. The idea behind Algorithm 1 is that $G(L)$ is connected during a period $[x, y]$ if and only if $G(L)$ has a spanning tree for every time instance in $[x, y]$.

Initially, Algorithm 1 finds a spanning tree of the input network $G(L)$ at time x. If no such tree exists, then at time x the network is disconnected and the algorithm terminates. If a spanning tree T exists at time x, then T remains connected until one (or more) of its edges stop being available. Denote by b_1

the first moment in time at which T disconnects. T consists now of a number of connected components and, in fact, T is a forest (collection of trees). The algorithm checks whether there are edges of $G(L)$ that are available at time b_1, which can be added to T and re-connect it. More specifically, the algorithm finds an edge that is available at time b_1 and has endpoints in different connected components of T. The algorithm adds that edge to T and checks if this addition re-connects it. If not, then it looks for yet another edge that is available at time b_1 and has endpoints in different connected components of (the current) T. This process continues until T is re-connected or we cannot find any more edges of $G(L)$ that are available at time b_1 and have endpoints in different connected components of T. If at any step of the process there do not exist edges that can re-connect T, then the algorithm returns that the network is disconnected. However, if we can find appropriate edges to re-connect T, then we form another spanning tree of the network, available from time b_1 onwards, and the same procedure continues. The algorithm answers that the network is connected if we form a spanning tree, all the edges of which are available until the end of the period in question, namely until time y.

Running Time. The running time of Algorithm 1 depends on the number of times that the spanning tree changes during $[x, y]$. The spanning tree can only change when one or more edges stop being available, so the above number is in general upper bounded by the total number of intervals assigned to the edges of the network:

$$M = \sum_{e \in E} |L_e|$$

Initially, to find $E(x)$ we need to look at every edge $e \in E$ and decide if x is between the start and finish time of one of e's availability intervals. Performing a binary search on the ordered set of start times *and* the ordered set of finish times of e's availability intervals, we can decide if $e \in E(x)$ in time $O(\log |L_e|)$. So, to compute $E(x)$ and check if it induces a spanning tree, we need time:

$$M' = O(\sum_{e \in E} \log |L_e|)$$

Next, time $O(n \log n)$ is required to sort the edges in T, where n is the number of vertices in the network. Then, for every time T changes, we need time M' to find the new set of available edges at the time. We need time $O(n)$ to find the connected components that can be re-connected by the addition of an available edge at the time and update T. Since we add at most $O(n)$ edges to re-connect T, the addition of all edges and the updates of T take a total of $O(n^2)$ time. Also, time $O(n)$ is required to insert the added edges in the sorted list A. Therefore, the running time of Algorithm 1 is $O\big(M' + n \log n + M \cdot (M' + n^2)\big) = O\big(M \cdot (M' + n^2)\big)$.

4   Large Connected Components During a Given Time Period

In this section, we examine if, given an interval temporal network $G(L)$ of n vertices, numbers $x, y \in \mathbb{R}^+$ and a parameter $0 \leq \varepsilon \leq 1$, we can find one or more *large enough* subsets of the vertices of G which remain connected and remain large within the time interval $[x, y]$. The matter of how large we want the components to be is handled by adjusting ε, which gives us a lower bound of $\varepsilon \cdot n$ on the size of the components we are looking for. In this section, we provide an algorithm that efficiently solves the above problem. Henceforth, a "large enough" connected component will be a component of size at least $\varepsilon \cdot n$.

Notice that any connected component C of $G(L)$, at time $t = x$, that is not large enough can be omitted by any algorithm that solves the above problem. Even if the vertices of C connect with more vertices in $G(L)$ at a later moment in time within $[x, y]$, resulting in a large enough connected component C' of $G(L)$ at that time, C' is not a component that was connected throughout $[x, y]$.

Theorem 2. *There is a polynomial-time algorithm which, given an interval temporal network $G(L)$ on n vertices and numbers $x, y \in \mathbb{R}^+$, $x < y$, returns all subgraphs of G of size $\varepsilon \cdot n$, $0 \leq \varepsilon \leq 1$, that remain connected and large (i.e., is always of size at least $\varepsilon \cdot n$) during the time period $[x, y]$. If $[x, a] \subseteq [x, y]$, $a \in [x, y]$, is the maximal sub-period of $[x, y]$ during which such a component remains connected, then the algorithm also returns the length of that period, $a - x$.*

Description of the Algorithm. Algorithm 2 receives as input an interval temporal network of n vertices and an interval $[x, y]$ during which we want to check whether one or more large components of the network remain connected. The algorithm also takes a non-negative parameter ε no larger than 1. This parameter defines how large we want our components to be; more specifically, the algorithm will only look for components of size (number of vertices) at least $\varepsilon \cdot n$. The algorithm returns all those subsets of the vertices of the initial graph, if any, that remain connected (and large) during $[x, y]$. Furthermore, it returns the duration of connectivity (*survival duration*) of any large enough component that was connected at time $t = x$ but disconnects at some point in $[x, y]$.

To do so, the algorithm initially checks which connected components, if any, are large enough at time x, and ignores all the rest. Then, the algorithm treats each and every one of these large components similarly, but separately. Namely, for each one of them the algorithm finds a spanning tree T and sorts all its edges according to the finish time of their availability interval, considering only the interval that includes time x. If the same tree remains connected during $[x, y]$, then the algorithm returns the respective component. Otherwise, if the tree disconnects at a moment t_0 in time, the algorithm employs a similar process to the one used in Algorithm 1, i.e., tries to reconnect the remainder of the tree via edges that are available at t_0. If T cannot be re-connected, then the algorithm

Algorithm 2. Connectivity of interval temporal graphs

Input: A temporal network $G(L)$ of n vertices, numbers $x, y \in \mathbb{R}^+$ such that $x < y$ and parameter $\varepsilon: 0 \leq \varepsilon \leq 1$

Output: All components of $G(L)$ of size $\varepsilon \cdot n$ that remain connected during the time interval $[x, y]$

1: Find the set $E(x)$ of available edges at time x, distinguish the connected components and delete those of size smaller than $\varepsilon \cdot n$;
2: **for** each of the remaining connected components **do**
3:   Find a spanning tree, T;
4:   $n' = |V(T)|$;
5:   Sort the edges in T according to the finish time of their availability interval;
     // For every edge in T, we only consider the interval that includes x
6:   Let $A = \{e_i, \text{ with interval } [a_i, b_i]: i = 1, \ldots, n-1\}$ be the sorted list;
7:   **if** $b_1 \geq y$ **then**
8:     **return**   V(T) **and** "Duration of survival of component $=$ " $y - x$;
9:   **else**
10:     $E' := \{e \in E(T) : b_e = b_1\}$;   // b_1 is the first time instance at which T becomes disconnected. E' is the set of edges of T that stop being available at time b_1.
11:     $T := T \setminus E'$;
12:     Remove E' from A;
13:     Let $T_1, T_2, \ldots, T_i, i \in \mathbb{N}$ be the connected components of T;
14:     **while** T is disconnected **and** $|V(T)| = n'$ **do**
15:       **if** $\exists j, k = 1, 2, \ldots, i : \exists e = (u, v) \in E(b_1) : u \in V(T_j) \wedge v \in V(T_k)$ **then**
16:         Find the T_j, T_k trees of T that e connects; //If there is an edge of G with endpoints in different connected components of T and is available at time b_1, then add it to T
17:         Merge T_j, T_k and e into a single tree;
18:         Update the number i of connected components of T;
19:         Insert e in the sorted list A;
20:       **else**
21:         **for** each connected component C of T with size smaller than $\varepsilon \cdot n$ **do**
22:           $T = T \setminus C$;
23:           **return** "Duration of survival of component $=$ " $b_1 - x$;
24:       **for** each connected component, C', of T **do**
25:         $T := C'$;
26:         $n' = |V(T)|$;
27:         Go to line 7;

---

checks the sizes of its connected components; it ignores those that are not large enough, while "processing" the rest similarly and separately as before. For each component that is ignored in the process, the algorithm returns the duration of its survival, meaning how long its vertices stayed connected since time x. The algorithm stops when there are no more components that are large enough or when the last component stays connected until time y.

Running Time. It is easy to see that the running time of the algorithm for each separate component is the same as the running time of Algorithm 1. Since there are at most $\frac{1}{\varepsilon}$ connected components of size at least $\varepsilon \cdot n$ in $G(L)$ during $[x, y]$, the running time of Algorithm 2 is $O\left(\frac{1}{\varepsilon}\left(M \cdot (M' + n^2)\right)\right)$.

5   Low Cost Maintenance of a Tree Structrure

In this section, we consider an interval temporal network on an underlying clique of n nodes, i.e., all $\binom{n}{2}$ links between nodes of the network *may* exist.

The connectivity of the network needs to be maintained at all moments in time via a tree structure, i.e., a spanning tree of the clique. Each node of the tree performs an individual application determined by the operator of the structure and each link (edge) is active (alive) during a time-interval also decided by the operator, after which the link fails. We have the liberty to provide the operator with extra edges from the clique to re-connect a spanning tree when a link fails; note here that after a new edge is added to the tree structure, the operator then assigns to it a "lifetime" interval, which is determined by the application, anew. The extra edges that we can provide come from the edges of the clique that are not currently used in the tree structure, i.e., a total of $\binom{n}{2} - (n-1)$ edges. We need to assign to every such edge e out of the $\binom{n}{2} - (n-1)$ an availability interval, I_e, so that when the tree structure becomes dis-connected, there is an appropriate such edge available to re-connect it. We call those edges *reserved* edges and the set that consist exactly of all those edges (with their availability intervals) *reservoir*, denoted by R.[3]

Definition 3 (Cost of the Reservoir). *The cost of the reservoir is defined as the sum, over all reserved edges, of the length of the edges' availability interval:*

$$c = \sum_{e \in R} |I_e|$$

Let T be the tree structure that is handled by the operator. We consider the time period between 0 and n and we assume that the breaks/failures in the connectivity of T happen once inside every consecutive time interval of length $\Delta \geq \alpha \log n$, for some $\alpha > 1.2$ (*Low-frequency-of-link-breaks assumption*). We are not able to predict when exactly the failures happen, nor are we able to foresee which link will fail next. We also assume worst case breaks in the tree topology within each Δ-interval. The trivial design of the availabilities of the reserved edges would be to make them all available throughout the considered time period $[0, n]$. However, this yields cost $c = \sum_{i=1}^{\binom{n}{2}-n+1} n \in O(n^3)$. We will

---

[3] Notice, here, the distinction between the *availability* of an edge and the *lifetime* of an edge: availability refers to the interval that we assign to a reserved edge with the purpose to re-connect the tree when it breaks, and lifetime refers to the interval that the operator assigns to an edge after it is inserted in the tree structure and is the time interval after which the respective link in the tree structure will fail.

show how to provide the network with available reserved edges with lower cost, so that the network connectivity is maintained with high probability (whp).[4] In order to re-connect the tree in the worst case of breaks in the tree topology, each reserved edge needs to have been randomly assigned to an availability interval to allow for the same probability of re-connection for all edges.

Theorem 3. *Let $\alpha \in \{x \in \mathbb{R} | x \geq 0.75\}$. If failures of the edges happen once in every consecutive $\Delta \geq \alpha \log n$ time-intervals, then there exists a reservoir of cost $O(n^2 \log n)$ that keeps a spanning tree available during $[0, n]$ whp.*

Proof. Partition the time interval $[0, n]$ into consecutive equisized sub-intervals $b_1, b_2, \ldots, b_{\frac{n}{\beta \log n}}$ of length $\beta \log n$, $\beta \in \mathbb{R}$, $0.75 \leq \beta \leq \alpha$, called *boxes*.[5] For every reserved edge $e \in R$ independently, select a box uniformly at random to be the availability interval of e. For every edge $e \in R$, the probability that e is assigned a particular box $b_i, i = 1, 2, \ldots, \frac{n}{\beta \log n}$ as its availability interval is:

$$Pr[I_e = b_i] = \frac{\beta \log n}{n}$$

Denote by m' the number of edges in R that are assigned to a particular box b_i, $i = 1, 2, \ldots, \frac{n}{\beta \log n}$, $m' = |\{e \in R : I_e = b_i\}|$. The expected value of m' is:

$$\mu = E[m'] = \frac{\beta \log n}{n} \cdot \left(\frac{n(n-1)}{2} - n + 1\right) = \frac{\beta n \log n}{2} - \frac{3\beta \log n}{2} + \frac{\beta \log n}{n}$$

By Chernoff bounds, we get that the probability that m' is *close* to the expected number of edges in a particular box $b_i, i = 1, 2, \ldots, \frac{n}{\beta \log n}$ is:

$$Pr[m' \in (1 \pm \frac{1}{2})\mu] \geq 1 - e^{-\frac{1}{4}\mu}$$

$$= 1 - e^{-\frac{1}{4} \cdot \frac{\beta \log n}{n} \cdot (\frac{n^2}{2} - \frac{n}{2} - n + 1)}$$

$$\geq 1 - \frac{1}{n^{\frac{\beta n}{16}}}, \text{ for } n \text{ large enough } (n \geq 6)$$

We now show that when a failure happens in T, we can whp find an edge in R which is available at that particular moment in time. Consider the specific box b_i that includes the time moment at which the failure in T happens. The number of edges in b_i that can re-connect T depends on where the failure happens, i.e., on the sizes[6] of the two connected components after the failure. If n_1 and n_2 are the sizes of the connected components of T after a failure, then the probability that a particular edge $e \in b_i$ can re-connect T after being added to the structure is:

---

[4] An event occurs with high probability if, for any $\gamma \geq 1$, the event occurs with probability at least $1 - \frac{c_\gamma}{n^\gamma}$, where c_γ depends only on γ.

[5] The last box is not necessarily of size exactly $\beta \log n$ but this does not affect the analysis.

[6] The size of a component is the number of its vertices.

$$Pr[e \in b_i \text{ re-connects } T] = \frac{n_1 \cdot n_2}{\frac{(n_1+n_2)\cdot(n_1+n_2-1)}{2}} \geq \frac{2n_1 n_2}{(n_1+n_2)^2}$$

The probability that no edge of b_i reconnects T after a failure is:

$$Pr[\text{no } e \in b_i \text{ re-connects } T] \leq \left(1 - \frac{2n_1 n_2}{(n_1+n_2)^2}\right)^{m'}$$

So, the probability that there is an edge in b_i that re-connects T is:

$$Pr[b_i \text{ re-connects } T] \geq 1 - \left(1 - \frac{2n_1 n_2}{(n_1+n_2)^2}\right)^{m'} \geq 1 - \left(1 - \frac{2n_1 n_2}{(n_1+n_2)^2}\right)^{\frac{3\mu}{2}}$$

In the worst case, T dis-connects into a component of size $n-1$ and a single vertex. So, we can reconnect T after a failure with probability:

$$Pr[b_i \text{ re-connects } T] \geq 1 - \left(1 - \frac{2(n-1)}{n^2}\right)^{\frac{3\beta n \log n}{4} - \frac{9\beta \log n}{4} + \frac{3\beta \log n}{2n}}$$

$$\geq 1 - \left(1 - \frac{2(n-1)}{n^2}\right)^{\frac{3}{4}\beta n \log n}$$

$$\geq 1 - \frac{1}{n^{0.9\frac{3}{2}\beta}}, \text{ for } n \text{ large enough } (n \geq 10)$$

$$= 1 - \frac{1}{n^{1.35\beta}} \xrightarrow{n \to +\infty} 1$$

We require $\beta \geq 0.75$ so that the above event happens whp. The probability that within the time period $[0, n]$, there is a box that will not re-connect T is:

$$Pr[\exists b_i, \ i = 1, \ldots, \frac{n}{\beta \log n} : b_i \text{ doesn't re-connect } T] \leq \sum_{i=1}^{\frac{n}{\beta \log n}} \frac{1}{n^{1.35\beta}}$$

$$= \frac{n}{\beta \log n} \cdot \frac{1}{n^{1.35\beta}} \xrightarrow{n \to +\infty} 0$$

So, we can almost surely[7] re-connect T during $[0, n]$ by employing the above random assignment of availability intervals to the reserved edges, having total cost $c = \sum_{i=1}^{\binom{n}{2}-n+1} \beta \log n \in O(n^2 \log n)$.   □

Conjecture. *If failures of the edges happen once in every consecutive $\Delta \geq \alpha \log n$ time-intervals, we conjecture that there is no reservoir of cost $o(n^2 \log n)$ that keeps a spanning tree available during $[0, n]$ whp.*

Open Problem 1. *For spanning tree breaks of frequency $o(\log n)$ within the time period $[0, n]$, the reservoir of Theorem 3 does not re-connect T whp. It remains an open question to derive a scheme that does so for breaks of so high frequency.*

Open Problem 2. *What is a low cost reservoir to maintain a spanning tree of the clique network, if the failures in the links of the tree happen randomly, e.g., if each link receives a lifetime given by the Exponential Distribution?*

---

[7] Note that increasing the size of the boxes by a constant factor, i.e., increasing the lower bound for β and α, can enforce the re-connection probability to also increase.

References

1. Akrida, E.C., Gąsieniec, L., Mertzios, G.B., Spirakis, P.G.: Ephemeral networks with random availability of links: diameter and connectivity. In: Proceedings of the 26th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA) (2014)
2. Angluin, D., Aspnes, J., Diamadi, Z., Fischer, M., Peralta, R.: Computation in networks of passively mobile finite-state sensors. Distrib. Comput. **18**(4), 235–253 (2006)
3. Avin, C., Koucký, M., Lotker, Z.: How to explore a fast-changing world (cover time of a simple random walk on evolving graphs). In: Aceto, L., Damgård, I., Goldberg, L.A., Halldórsson, M.M., Ingólfsdóttir, A., Walukiewicz, I. (eds.) ICALP 2008, Part I. LNCS, vol. 5125, pp. 121–132. Springer, Heidelberg (2008)
4. Bui-Xuan, B.-M., Ferreira, A., Jarry, A.: Computing shortest, fastest, and foremost journeys in dynamic networks. Int. J. Found. Comput. Sci. **14**(2), 267–285 (2003)
5. Casteigts, A., Flocchini, P., Quattrociocchi, W., Santoro, N.: Time-varying graphs and dynamic networks. Int. J. Parallel Emergent Distrib. Syst. (IJPEDS) **27**(5), 387–408 (2012)
6. Clementi, A.E.F., Macci, C., Monti, A., Pasquale, F., Silvestri, R.: Flooding time of edge-markovian evolving graphs. SIAM J. Discrete Math. (SIDMA) **24**(4), 1694–1712 (2010)
7. Cooper, C., Klasing, R., Radzik, T.: A randomized algorithm for the joining protocol in dynamic distributed networks. Theor. Comput. Sci. **406**(3), 248–262 (2008)
8. Duchon, P., Duvignau, R.: Local update algorithms for random graphs. In: Pardo, A., Viola, A. (eds.) LATIN 2014. LNCS, vol. 8392, pp. 367–378. Springer, Heidelberg (2014)
9. Dutta, C., Pandurangan, G., Rajaraman, R., Sun, Z., Viola, E.: On the complexity of information spreading in dynamic networks. In: Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 717–736 (2013)
10. Fleischer, L., Skutella, M.: Quickest flows over time. SIAM J. Comput. **36**(6), 1600–1630 (2007)
11. Fleischer, L., Tardos, É.: Efficient continuous-time dynamic network flow algorithms. Oper. Res. Lett. **23**(3–5), 71–80 (1998)
12. Gavoille, C., Peleg, D., Perennes, S., Raz, R.: Distance labeling in graphs. In: Proceedings of the 12th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), pp. 210–219 (2001)
13. Katz, M., Katz, N.A., Korman, A., Peleg, D.: Labeling schemes for flow and connectivity. SIAM J. Comput. **34**(1), 23–40 (2004)
14. Kempe, D., Kleinberg, J.M., Kumar, A.: Connectivity and inference problems for temporal networks. In: Proceedings of the 32nd Annual ACM Symposium on Theory of Computing (STOC), pp. 504–513 (2000)
15. Koch, R., Nasrabadi, E., Skutella, M.: Continuous and discrete flows over time - a general model based on measure theory. Math. Methods OR **73**(3), 301–337 (2011)
16. Kuhn, F., Lynch, N.A., Oshman, R.: Distributed computation in dynamic networks. In: Proceedings of the 42nd Annual ACM Symposium on Theory of Computing (STOC), pp. 513–522 (2010)
17. Mertzios, G.B., Michail, O., Chatzigiannakis, I., Spirakis, P.G.: Temporal network optimization subject to connectivity constraints. In: Fomin, F.V., Freivalds, R., Kwiatkowska, M., Peleg, D. (eds.) ICALP 2013, Part II. LNCS, vol. 7966, pp. 657–668. Springer, Heidelberg (2013)

18. Michail, O., Chatzigiannakis, I., Spirakis, P.G.: Causality, influence, and compu-
    tation in possibly disconnected synchronous dynamic networks. In: Baldoni, R.,
    Flocchini, P., Binoy, R. (eds.) OPODIS 2012. LNCS, vol. 7702, pp. 269–283.
    Springer, Heidelberg (2012)
19. Molloy, M., Reed, B.: Graph Colouring and the Probabilistic Method. Algorithms
    and Combinatorics, vol. 23. Springer, Heidelberg (2002)
20. O'Dell, R., Wattenhofer, R.: Information dissemination in highly dynamic graphs.
    In: Proceedings of the 2005 Joint Workshop on Foundations of Mobile Computing
    (DIALM-POMC), pp. 104–110 (2005)
21. Scheideler, C.: Models and techniques for communication in dynamic networks. In:
    Alt, H., Ferreira, A. (eds.) STACS 2002. LNCS, vol. 2285, pp. 27–49. Springer,
    Heidelberg (2002)